High Performance Network Simulation using OPNET

Sanjay Hegde, Vikram Mallikarjuna, Cynthia Hood
Illinois Institute of Technology
Chicago, IL 60616
E-mail: { hegdsan, mallvik, hood}@iit.edu

Abstract

As our reliance on high-performance systems increases, it is increasingly important to understand the behavior of these systems. Our research focuses on this understanding in the context of fault and performance management. Toward this goal, we are studying the Access Grid from the measurements taken in the actual system and from an OPNET simulation. This paper describes the network piece of the simulation. The Abilene high performance backbone and GigaPoPs are described. Issues to be addressed are discussed.

Introduction

The advanced technologies used in both research and education are increasingly based on high performance networks that provide high-speed access to geographically distributed data and resources [7]. Our research focuses on fault and performance management for these distributed systems. We are currently studying this problem in the context of the Access Grid (AG) [11]. The AG is a distributed collaboration environment that uses multicast technology over high-speed networks. Along with the study of the real system, we create a simulation of the AG. The simulation provides the opportunity for more extensive study of fault and performance problems and the implementation of new mechanisms.

The simulation of the AG can be broken into two pieces, the network and the AG node. In this paper, we focus on the network simulation. As part of the network simulation we simulate the Abilene backbone. Abilene is a high-performance network developed by the University Corporation for Advanced Internet Development (UCAID) in partnership with Cisco Systems, Juniper Networks, Nortel Networks, Qwest Communications, and Indiana University. Abilene provides a backbone network for the Internet2 community. As shown in Figure 1, it consists of high-performance IP routers. These routers are accessible to regional aggregation points (i.e.GigaPoPs) in several dozen locations nationwide [12].

High Performance Network Overview

Abilene is an advanced backbone network that provides high-speed, low latency, best effort, nationwide connectivity to support the development of advanced Internet applications [8]. It is a packet-over-SONET (POS) network made up of router nodes. Each Router nodes consists of a Cisco 12000 series GSR router, a Unix based computer for measurements and network management, and a Cisco 3640 remote access router to support out-of-band access from the Abilene NOC (Network Operations Center). The router nodes are connected via OC3 (155Mbps), OC12 (622Mbps) and OC48 (2.4 Gbps) SONET links.

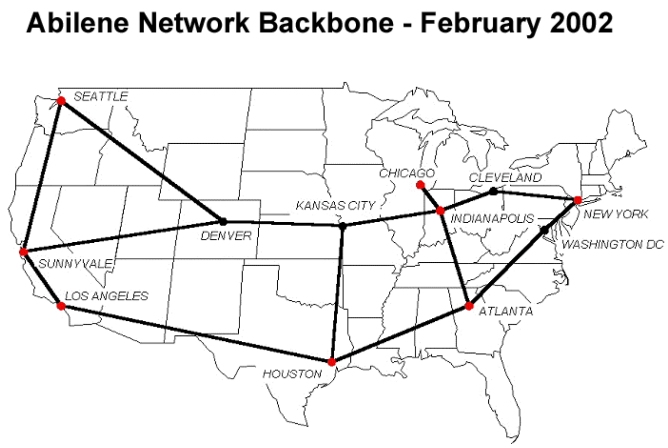

Figure 1: Abilene Network (Courtesy: http://www.ucaid.org/abilene/html/maps.html)

The partners mentioned above operate and manage the Abilene network. Abilene inter-connects other national and international high-performance networks. Non-US National Research Educational Networks (NERNs) are connected to Abilene via an International Interconnection Point. Currently more than fifty-five peer networks are connected to Abilene. Some of the national peer networks are the Defense Research and Engineering Network (DREN), Esnet(Energy Sciences Network), and NISN(NASA Integrated Services Network.). International peer networks include the European Organization for Nuclear Research (CERN), and Australia's Academic and Research Network (AARNet). The interconnection of different peer networks and participants are shown in Figure 2.

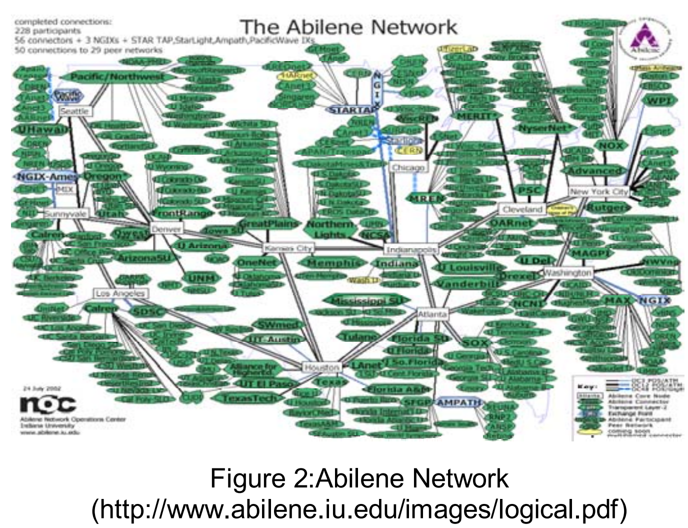

Figure 2:Abilene Network (http://www.abilene.iu.edu/images/logical.pdf)

Internet2 participants may connect to Abilene through a GigaPoP, a direct connection to a core router, or through other participants. As of October 2001, there were 51 direct connections to Abilene.

GigaPoP Overview

A GigaPoP is an aggregate point of presence (POP) with gigabit capacity [4]. Traffic from various networks is aggregated at the GigaPoP and then routed to Abilene. This hierarchical model facilitates the shared use of distributed resources and minimizes the manual effort necessary to manage the network. Figure 3 shows the basic structure of the GigaPoP. The Internet2 engineering committee has made the following recommendations for a GigaPoP site.

- A layer 2 GigaPoP should support IPV6 along with IPV4.
- All the usual supporting protocols are assumed to be available along with the IGMP and RSVP to support multicasting and resource reservation.
- Packet loss within the GigaPoP is close to zero.
- One or more very high-capacity advanced function packet data switch/routers capable of supporting OC12 links.
- Although it is an interconnection point between different wide area network and the backbone, the GigaPoP should not become a bottleneck in access to wide area communications services.
- The GigaPoP internal design must be able to handle the aggregate throughput demanded by all local participants and wide area connections.
- A part of the GigaPoP architecture must be integrated data collection with appropriate safeguards but with enough detail and accuracy to support serious study and analysis.

Currently there are approximately 30 GigaPoPs [1] connected to Abilene. These GigaPoPs service more than 70% of the participants. On average there will be 8-10 networks connected to a GigaPoP. The number depends on the interface provided by the ATM switch and router. A common component of the GigaPoPs is an ATM switch and multicast enabled router.

GigaPoPs are divided into two broad types:
Type I GigaPoPs are simple in architecture and serve only Internet 2 (I2) members. These GigaPoPs will route the traffic to other GigaPoPs through one or two connection hops.
Type II GigaPoPs are relatively complex. They serve both I2 members and other I2 member connected networks. Type II GigaPoPs also provide QOS and enhanced security.

External connections to GigaPoP ATM switching elements may be direct SONET circuits from campus ATM switches or other GigaPoPs. The connections may also be full ATM service from commercial carriers. The ATM switching elements multiplex the link level bandwidth through permanent or switched virtual circuits (PVCs or SVCs). In this way, the intra- and inter-GigaPoP connectivity can be optimized and separate bandwidth can be allocated to test bed or other special purpose requirements [1].

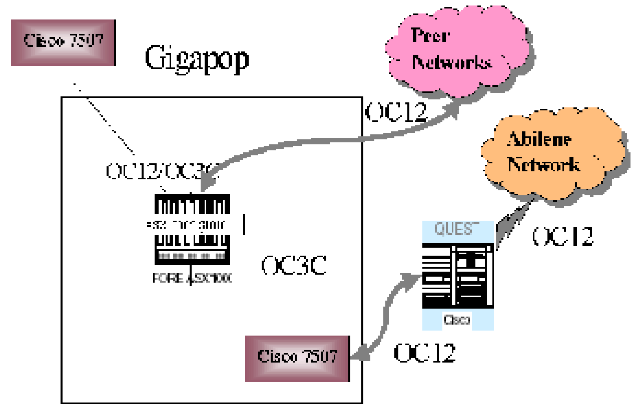

Figure 3:GigaPoP Model

Services provided by the Abilene

The services provided by the Abilene network is to support bandwidth-intensive applications such as distributed computing, real-time access and control of remote instrumentation, and efficient multicasting to support video and advanced multimedia collaboration such as Access Grid.

Access Grid Model

The AG is the aggregation of resources that can be used to support human interaction across the grid. It consists of multimedia display, presentation and interactions environments, interfaces to grid middleware, interfaces to visualization environments. The AG supports large-scale distributed meetings, collaborative work sessions, seminars, lectures, tutorials and training [11]. An AG node consists of four computers. A display computer receives the video streams from the grid and projects to a screen using projectors. The display/capture computer captures the video information and sends it to the grid. The audio capture computer sends and receives the audio information to and from the grid. The control system is used to control the audio echo cancellation device. The basic AG node configuration is shown in Figure 4.

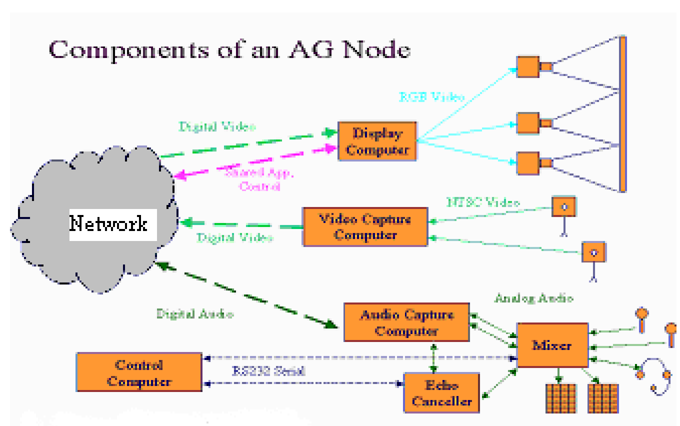

Figure 4: Components of AG node (This picture has been taken from ANL website and modified)

GigaPoP OPNET model

Each GigaPoP consists of at least one ASX1000 or LS1010 ATM switch and a router. OPNET provides models for the Cisco Light Stream 1010 switch so this is used. Since OPNET does not provide ASX1000 models, the LS1010 is used in its place. The provided multicast enabled Cisco 12000 series router and 7000 series routers are used to build the GigaPoPs. Information about the configuration and interfaces of the Abilene routers can be obtained from the Abilene website [8]. Information about the interconnection of other networks to Abilene can be obtained from the website for the particular GigaPoP. The GigaPoP and the core node router components constructed in this project are as shown in Figure 5.

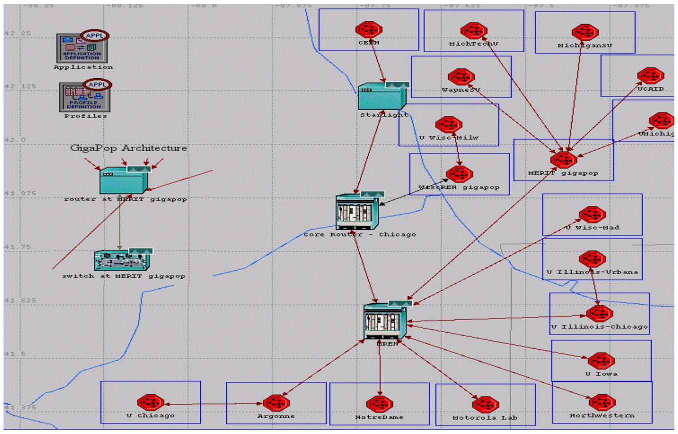

Figure 5: GigaPoP and Core node OPNET model

OPNET Access Grid Model

The OPNET AG model (Figure 6) consists of duel processor, 866MHz, PIII Dell Precision 420 Workstations for Display machine and Video machine. Audio and Control Machines are 700MHz, PIII, Dell Precision 420 workstations. These four systems connected to advanced 16 port Ethernet hub and the hub will be connected to the Cisco 7000 series router via 100MBps link.

Testing the Simulation Model

The first step toward validating the OPNET simulation is to ensure that the models are working as expected. Since this will ultimately be a very large simulation, we started with the local piece of the network. We began by simulating the scenario shown in Figure 5. This represents of basic connection of one Chicago core node interconnection network.

We ran the simulation on a Sun Ultra-5 Spark with 256 MB RAM and Sun OS Release 5.8. OPNET version 8.1 was used for the simulation. The network consists of 153 Cisco routers, 5 ATM switches, 360 Ethernet Advanced workstations and 120 hubs. The application used for the simulation was low resolution video. All the simulations are run for 5 minutes of simulated time with different seeds and loads. The performance measurements are logged. The metric we use includes traffic between the links(bits/sec), input and output packets in error, packets discarded by router and the delay. Measurements show that we were able to simulate a traffic rate of approximately 3Mbps on the 1Gbps link (see Figure 8) with the single video conferencing application. The end-to-end delay (Figure 7) for the video conferencing is approximately 3.5ms. The measured packet loss ratio was close to zero (Figure 9). Also most of the GigaPoPs we constructed were able to handle the traffic sent by its connectors without any performance slowdown. The next steps are to increase the network size and to validate the simulation against data collected from the AG.

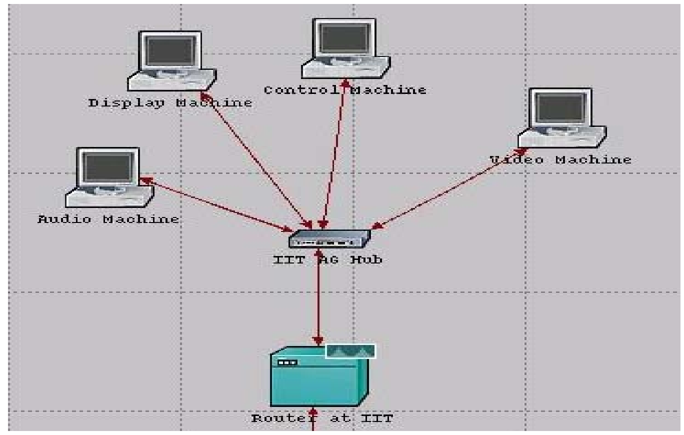

Figure 6: Access Grid OPNET Model

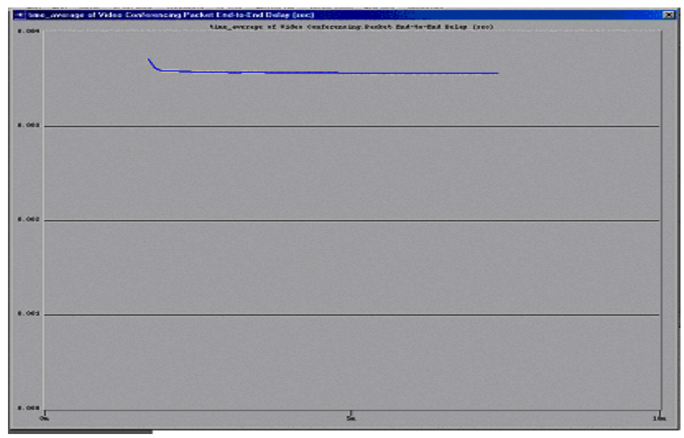

Figure 7: End to End Delay

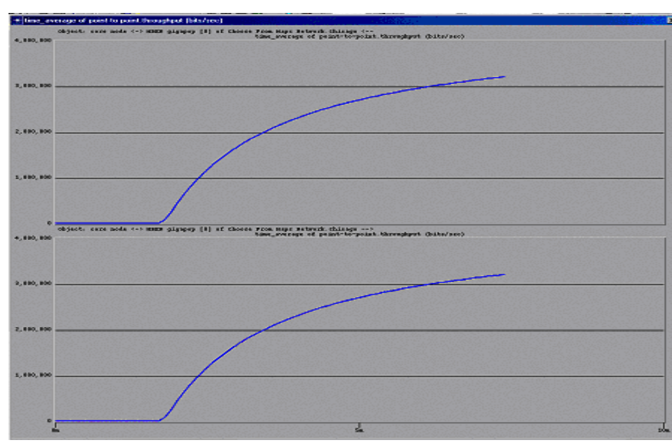

Figure 8: Throughput Between the Router and GigaPoP

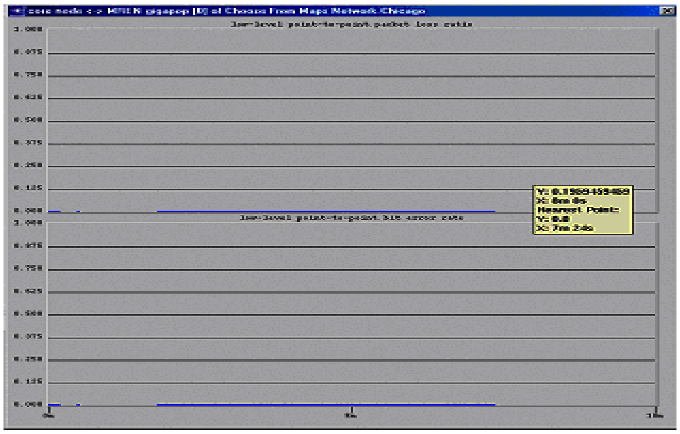

Figure 9: Packet Loss Ratio

Conclusion

The collection of models provided by OPNET allows simulation models of high-performance networks to be created quickly. In our research, we are using the AG as a testbed. The AG simulation described above will allow us to better understand the behavior of the AG under stress and to implement mechanisms for fault and performance management.

This effort is in the early stages and the future work is in two key directions. The first direction is to validate the simulation with data from collected from the live AG. This will help us make the simulation as close to reality as possible and will allow us to understand the limitations of the simulation. The second direction is scalability of the simulation. Simulating a large-scale high-performance network takes significant resources and time. We are looking at ways to modify the simulation to make it more efficient.

The support of NSF CAREER 9984811 is gratefully acknowledged.

References

[1] GigaPoPs structure and services. http://www.internet2.edu/ html/GigaPoPs.html

[2] Internet2 GigaPoPs. http://www.internet2.edu/ resources/ Internet2GigaPoPsMap.PDF

[3] Internet2 GigaPoP workshop, April 1997 http://www. internet2. edu/html/GigaPoP_workshop.html

[4] Internet2 and GigaPoP development white paper, 1997. http://www.cisco.com/warp/public/750/aii/i2_GigaPoP.pdf

[7]http://www.indiana.edu/~ovpit/presentations/cenic2000/sld002.htm

[8] Abilene web site. http://www.ucaid.org/abilene/

[9] Abilene Premium Service Test Program Draft, 2000. http://www.internet2.edu/abilene/qos/aps.pdf

[10] Internet2 and GigaPoP Development White Paper, 1997. http://www.cisco.com/warp/public/750/aii/i2_GigaPoP.pdf

[11] Access Grid web site. www.accessgrid.org

[12] Frequently asked questions http://www.internet2.edu/ Abilene / html/faq-general.html